

The Wilson Machine for Image Modeling

Saeed Saremi¹ & Terrence J. Sejnowski¹

¹*Salk Institute, 10010 N Torrey Pines Rd, La Jolla, CA 92037.*

Learning the distribution of natural images is one of the hardest and most important problems in machine learning, computer vision, and artificial intelligence¹⁻³. Intelligent agents must navigate in the natural world; having the right model of natural images is thus essential for their success in visual tasks^{4,5}. The problem remains open, however, because the enormous complexity of the structures in natural images spans all length scales⁵⁻⁷. Here we break down the complexity of the problem and show that the hierarchy of structures in natural images fuels a new class of learning algorithms based on the theory of critical phenomena and stochastic processes. We approach this problem from the perspective of the theory of critical phenomena, which was developed in condensed matter physics to address problems with infinite length-scale fluctuations⁷⁻⁹, and build a framework to integrate the criticality of natural images¹⁰⁻¹³ into a learning algorithm. The problem is broken down by mapping images into a hierarchy of binary images, called bitplanes.

In this representation, the top bitplane is critical, having fluctuations in structures over a vast range of scales. The bitplanes below go through a gradual stochastic heating process to disorder. We turn this representation into a directed probabilistic graphical model¹⁴, transforming the learning problem into the unsupervised learning of the distribution of the critical bitplane and the supervised learning of the conditional distributions for the remaining bitplanes. We learnt the conditional distributions by logistic regression in a convolutional architecture. Conditioned on the critical binary image, this simple architecture can generate large, natural-looking images, with many shades of gray, without the use of hidden units, unprecedented in the studies of natural images. The framework presented here is a major step in bringing criticality and stochastic processes to machine learning and in studying natural image statistics.

Learning the distribution of natural images remains one of the hardest problems in unsupervised learning. Structures in natural images are complex, varied, and most importantly span a vast range of length scales, from the smallest within a few dozen pixels to large structures the size of the image itself. In this respect, they are very similar to critical points realized for physical systems near continuous (second-order) phase transitions⁷. There is a deep connection underlying this

similarity as critical large-scale fluctuations emerge after mapping natural images to a stack of binary images^{10–12}. Here we introduce an algorithm to utilize this criticality in learning the distribution of natural images.

The binary representation of Ref. 10 contained ordered and disordered phases, with a critical bitplane close to a phase transition. The critical bitplane underlies the large-scale fluctuations of the structures in natural images. Here we build on those results and exploit the binary representation, turning it into a directed probabilistic graphical model¹⁴ by placing the critical bitplane at the root of the directed graph. This transforms learning the distribution of natural images into the *unsupervised* learning of the distribution of the critical bitplane and the *supervised* learning of the conditional distributions of all other bitplanes in a parent-child hierarchy outlined below. The network architecture for learning is simple. We approximate the conditional distributions by logistic regression, where the weights to the children nodes are learnt in a convolutional architecture by weight-sharing.

Turning unsupervised learning into supervised learning has a rich history and goes back to the wake-sleep algorithm and the Helmholtz machine, which co-trained inference (the wake phase) and generative models (the sleep phase) together^{15,16}. Here, in contrast, the supervision is given by the input itself. Instead

of learning deep architectures with hidden units^{1,2}, we show that the visible hierarchy in the form of critical structures and stochastic processes can be utilized for learning. These two types of hierarchies complement each other. In particular, the visible hierarchy can be modeled using hidden units in the deep learning framework.

The starting point is to construct an *exact* binary representation so that the top bitplane is critical and the bottom ones go through a stochastic heating process to disorder. In the binary representation that was studied in 10,12, the critical bitplane was in the middle of the hierarchy, with different types of structures emerging in relation to it, “ordered” towards the top, and “disordered” towards the bottom. It is not clear in that representation how the bitplanes just above and just below the critical bitplane could influence each other in a directed graph. However, in the new representation, the causal direction is clear: the stochastic heating process from criticality (top) to disorder (bottom).

The new representation is a variation on the original¹⁰ but it turns out to be key in how we map the binary representation to a directed graphical model. We first give a brief review of the original representation. The analog values I (assumed to be non-negative integers for simplicity) were mapped to a binary

representation $\{B_1, B_2, \dots, B_\Lambda\}$ by the following decomposition

$$I = \sum_{\lambda=1}^{\Lambda} 2^{\Lambda-\lambda} B_\lambda, \quad (1)$$

where $B_\lambda \in \{0, 1\}$ were found iteratively starting from $\lambda = 1$. The visual representation of the binary map from analog values to their corresponding bits is given in Fig. 1. The binary map is extended to any of array, mapping the gray-scale image \mathcal{I} to a stack of binary images $\mathbb{B} = \{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_\Lambda\}$.

The goal is to build a directed probabilistic graphical model on top of this representation, where

$$P(\mathbb{B}) = P(\mathcal{B}_1)P(\mathcal{B}_2|\mathcal{B}_1) \cdots P(\mathcal{B}_\Lambda|\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_{\Lambda-1}). \quad (2)$$

However, to find good estimates for the conditional distributions, the structure of the directed graph should be dictated by a “physical” causal process. In a simple tweak, we change the binary representation so that \mathcal{B}_1 is critical and the bitplanes below go through a gradual stochastic process to disorder. The stochastic process becomes the causal direction, and the parent-child hierarchy becomes clear: from criticality to disorder.

In the new representation, all the half-point dividing lines in Fig. 1 are changed to the median values of their corresponding intervals. This hierarchi-

cal median-thresholding is *exactly* equivalent to applying *half-point* thresholding of Fig. 1 after first performing histogram equalization

$$I \rightarrow I' = \int_{-\infty}^I P(x)dx, \quad (3)$$

where $P(x)$ is the marginal distribution of the pixel intensities. The half-point thresholding of the histogram-equalized image is equivalent to median-thresholding of the original image because histogram equalization does not change the rank of pixel values. In fact, any strictly monotonic transformation of \mathcal{I} yields the same binary representation \mathbb{B} under median-thresholding. The illustration of this new binary representation for an image in the Geisler database¹⁷ is given in Fig. 2.

The first bit in the new representation $\mathcal{B}_1 = \mathcal{I} > \text{median}(\mathcal{I})$ is the median-thresholded image, which was studied at length in 11–13. In particular, it was shown that \mathcal{B}_1 is at the onset of a percolation transition, with connected clusters governed by scaling laws on all length scales¹². In fact, these rich connected clusters at the top of the graphical model provide the fuel to the algorithm presented here.

What we gain in the new representation is what we planned to achieve in that, as λ increases, the bitplanes \mathcal{B}_λ go through a gradual stochastic heating pro-

cess to disorder (see Fig. 2), thus placing them in the parent-child hierarchy of the directed graphical model of Eq. 2. In summary, in the new representation, $\mathcal{B}_1 = \mathcal{I} > \text{median}(\mathcal{I})$ is critical, at the root of the graph, and \mathcal{B}_λ is the parent of $\mathcal{B}_{\lambda'}$ for $\lambda < \lambda'$. The stochastic process from criticality to disorder *is* the hierarchy, seen in Fig. 2 by the islands of connected clusters gradually dissolving in a sea of noise. With this construction of the graphical model, rooted in statistical physics and in the theory of critical phenomena, we find estimates for the conditional distributions $P(\mathcal{B}_\lambda | \mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_{\lambda-1})$.

The network architecture for learning the conditional distributions is very simple. We first assumed that pixels in the bitplane \mathcal{B}_λ are independent conditioned on the bitplanes above. This might appear as a strong assumption for bitplanes close to \mathcal{B}_1 but since the critical fluctuations happen at *infinite* scales, large structures in \mathcal{B}_1 induce marginal interactions in the bitplanes below. We estimated the conditional distributions $P(\mathcal{B}_\lambda | \mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_{\lambda-1})$ by taking $L \times L$ patches and performing logistic regression for bitplane \mathcal{B}_λ taking bitplanes $\{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_{\lambda-1}\}$ as input. Since natural images are translation-invariant, we assumed weight-sharing, making the logistic regression convolutional.

We performed the learning of the convolutional weights by maximizing the

log likelihood with a single batch using the second-order Newton’s method. The batch contained 10^5 samples of size 41×41 , taken from 1024 images of size 2844×4284 pixels from the Geisler database of natural images¹⁷. The learnt center-surround receptive field for $P(\mathcal{B}_2|\mathcal{B}_1)$ is shown in Fig. 3(a). Receptive fields for other conditional distributions have the same center-surround form. The results reported here were obtained by having ℓ_2 prior on the weights; results were qualitatively similar, however, after setting the regularization term to 0.

After learning the convolutional weights, conditioning on the first bitplane, we found the logistic activations of the lower bitplanes, up to \mathcal{B}_8 . A key characteristic of natural images is the hierarchy in object sizes. This appears as connected clusters in the bitplane representation¹². The peaks in Fig. 3(b) is the signature of those large clusters. To preserve the bigger clusters in lower bitplanes, which are key in natural images, we adopted a winner-take-all strategy, except for the interval $[0.4, 0.6]$. In that interval we took samples according to the logistic activation. With this algorithm, we can easily generate binary images by the increasing order in λ , dictated by the structure of the directed graph. For example, the image in Fig. 4 was generated by conditioning on the bitplane \mathcal{B}_1 of Fig. 2, generating bitplanes \mathcal{B}_2 up to \mathcal{B}_8 with logistic activations, and combining $\{\mathcal{B}_1, \mathcal{B}_2, \dots, \mathcal{B}_8\}$ to

obtain the gray-scale image.

The model performance was evaluated by the normalized mean square error and the conditional log likelihood scores. The normalized mean square error for the generated images, averaged over the database is 0.0588. The same error measured for the null model, where the convolutional weights/biases are set to zero, is 0.1244. For the density estimation measure, the average negative log likelihood for the conditional probabilities were evaluated and are given in Table 1. Note that in this framework, the log likelihood for all bitplanes \mathbb{B} combined is *not* a good measure to compare models, in part because the analog image is obtained by a *weighted* sum of the bitplanes. In other words, it is (much) more important to have a high log likelihood score for bitplanes closer to the critical bitplane \mathcal{B}_1 than the ones further below.

The state of the art in natural image modeling has been limited to texture synthesis, dead leaf images, small patches, and large images that do not look natural or realistic — see Ref. 3 for the most recent review. That being said, the direct comparison cannot be made to other models at the present stage, since the results here were obtained by conditioning on the critical bitplane \mathcal{B}_1 with its rich structures. Our results, however, point to a new research direction in image mod-

eling, as learning the prior on the critical bitplane \mathcal{B}_1 will transform this algorithm to a fully probabilistic model and a very powerful model for natural images.

This novel framework is named after *Kenneth G Wilson (1936 – 2013)* for his immense contributions to our understanding of the nature of criticality. Here, the criticality of natural images plays the central role in the learning algorithm, and in contrast with the research focus in the deep learning community on learning deep architectures with hidden units, we showed that the hierarchy of *visible* structures by itself fuels learning in the Wilson machine.

Along these lines, there has been a recent work in learning data distributions by turning them into noise through stochastic processes, with the use of hidden units¹⁸. Learning in that framework is to figure out how to reverse the stochastic process, thus going from noise to data. In contrast, in the Wilson machine described here, the learning happens due to the stochastic process that is already present in the data, and is part of the data itself. In addition, the results here point to a new direction in image compression as learning a better probabilistic model in this framework might push the (lossy) compression limit to 1 bit/pixel (i.e. the first bitplane \mathcal{B}_1 .) Finally, even though we focused on natural images, the framework here is general and should be especially powerful for analog signals with

very long correlation lengths/times.

1. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
2. Schmidhuber, J. Deep learning in neural networks: An overview. *Neural Networks* **61**, 85–117 (2015).
3. Gerhard, H. E., Theis, L. & Bethge, M. Modeling natural image statistics. In *Biologically-inspired Computer Vision: Fundamentals and Applications* (Wiley VCH, 2015).
4. Marr, D. *Vision: A computational investigation into the human representation and processing of visual information* (MIT Press, 2010).
5. Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annual review of neuroscience* **24**, 1193–1216 (2001).
6. Mandelbrot, B. B. *The fractal geometry of nature* (Macmillan, 1983).
7. Wilson, K. G. Problems in Physics with Many Scales of Length. *Scientific American* **241**, 140–157 (1979).

8. Wilson, K. G. & Kogut, J. The renormalization group and the ϵ expansion. *Physics Reports* **12**, 75–199 (1974).
9. Ma, S. *Modern theory of critical phenomena* (Westview Press, 2000).
10. Saremi, S. & Sejnowski, T. J. Hierarchical model of natural images and the origin of scale invariance. *Proceedings of the National Academy of Sciences* **110**, 3071–3076 (2013).
11. Saremi, S. & Sejnowski, T. J. On criticality in high-dimensional data. *Neural computation* **26**, 1329–1339 (2014).
12. Saremi, S. & Sejnowski, T. J. Correlated percolation, fractal structures, and scale-invariant distribution of clusters in natural images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2015).
13. Stephens, G. J., Mora, T., Tkačik, G. & Bialek, W. Statistical Thermodynamics of Natural Images. *Physical Review Letters* **110**, 018701 (2013).
14. Koller, D. & Friedman, N. *Probabilistic graphical models: principles and techniques* (MIT press, 2009).
15. Hinton, G. E., Dayan, P., Frey, B. J. & Neal, R. M. The “wake-sleep” algorithm for unsupervised neural networks. *Science* **268**, 1158–1161 (1995).

16. Dayan, P., Hinton, G. E., Neal, R. M. & Zemel, R. S. The Helmholtz machine. *Neural computation* **7**, 889–904 (1995).
17. Geisler, W. S. & Perry, J. S. Statistics for optimal point prediction in natural images. *Journal of Vision* **11**, 14 (2011).
18. Sohl-Dickstein, J., Weiss, E. A., Maheswaranathan, N. & Ganguli, S. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. *arXiv preprint arXiv:1503.03585* (2015).

Acknowledgements We acknowledge conversations with Ruslan Salakhutdinov, comments from Lyle Muller on the manuscript, and the support of The Howard Hughes Medical Institute.

Correspondence Correspondence and requests for materials should be addressed to S.S. (email: saeed@salk.edu).




Figure 1: **The binary representation of analog values.** The first three bits in the binary representation of the Λ -bit integer value I in the range $[0, 2^\Lambda - 1]$ is shown schematically. Each bit divides the interval in half iteratively, starting from the most significant bit B_1 .




Figure 2: **Hierarchical median-thresholding, the critical bitplane, and the stochastic heating process to disorder.** An image \mathcal{I} (2844×4284 pixels) in the Geisler database of natural images and its binary representation up to bitplane \mathcal{B}_7 are shown. The critical bitplane \mathcal{B}_1 goes through a gradual stochastic heating process to disorder.




Figure 3: **Learnt receptive fields and logistic activations.** **a**, The plot of the learnt receptive field for the conditional distribution $P(\mathcal{B}_2|\mathcal{B}_1)$ is given. The size of the receptive field is 41×41 pixels. The receptive fields to other bitplanes have the same center-surround form. **b**, The normalized histogram of the logistic activation function after applying the convolution filter in (a) to the bitplane \mathcal{B}_1 from Fig. 2 is plotted here. The peaks come from the large connected clusters in \mathcal{B}_1 .



Figure 4: **Image generation in the Wilson machine.** The image on top (2564×4004 pixels) was generated by conditioning on the bitplane \mathcal{B}_1 of Fig. 2 and obtaining other bitplanes up to \mathcal{B}_8 by the logistic activations from the learnt convolutional filters. The bitplanes were combined to obtain the gray-scale image. The area enclosed in the yellow rectangle is blown up on the left and the corresponding region in the original image is given on the right.

$\text{NLL}(\mathcal{B}_2 \mathcal{B}_1)$	$\text{NLL}(\mathcal{B}_3 \mathcal{B}_1, \mathcal{B}_2)$	$\text{NLL}(\mathcal{B}_4 \mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3)$
0.2369	0.2901	0.2749

Table 1: **Density estimation results.** The results on the density estimation of the conditional distributions are given. The average negative log likelihood (NLL) is reported in bits/pixel. There were about 10^{10} pixels in the experiments. The negative log likelihood of the null model, corresponding to white noise, is 1 bit/pixel.